arXiv:0908.0863v1 [astro-ph.IM] 6 Aug 2009

# OpenDDA: A Novel High-Performance Computational Framework for the Discrete Dipole Approximation

James Mc Donald* Aaron Golden* S. Gerard Jennings†

August 6, 2009

[1]Computational Astrophysics Laboratory, I.T. Department, National University of Ireland, Galway

[2]School of Physics & Environmental Change Institute, National University of Ireland, Galway

PROPOSED RUNNING HEAD: OPENDDA: HIGH-PERFORMANCE DDA

James Mc Donald[1]
Computational Astrophysics Laboratory, Department of Information Technology,
National University of Ireland, Galway
Email: james@it.nuigalway.ie Phone: 353 91 495632 Fax: 353 91 494584

Aaron Golden
Computational Astrophysics Laboratory, Department of Information Technology,
National University of Ireland, Galway
Email: agolden@it.nuigalway.ie Phone: 353 91 493549 Fax: 353 91 494501

S. Gerard Jennings
School of Physics & Environmental Change Institute,
National University of Ireland, Galway
Email: gerard.jennings@nuigalway.ie Phone: 353 91 492704 Fax: 353 91 495515

[1]Corresponding author: James Mc Donald (james@it.nuigalway.ie)

**Abstract**

This work presents a highly-optimised computational framework for the Discrete Dipole Approximation, a numerical method for calculating the optical properties associated with a target of arbitrary geometry that is widely used in atmospheric, astrophysical and industrial simulations. Core optimisations include the bit-fielding of integer data and iterative methods that complement a new Discrete Fourier Transform (DFT) kernel, which efficiently calculates the matrix-vector products required by these iterative solution schemes. The new kernel performs the requisite 3D DFTs as ensembles of 1D transforms, and by doing so, is able to reduce the number of constituent 1D transforms by 60% and the memory by over 80%. The optimisations also facilitate the use of parallel techniques to further enhance the performance. Complete OpenMP-based shared-memory and MPI-based distributed-memory implementations have been created to take full advantage of the various architectures. Several benchmarks of the new framework indicate extremely favourable performance and scalability.

Keywords: Discrete Dipole Approximation, optimisation, matrix-vector product, CG-FFT, parallel algorithm

# 1 Introduction

The Discrete Dipole Approximation (DDA) is an intuitive and extremely flexible numerical method for calculating the optical properties associated with a target of arbitrary geometry, whose dimensions are comparable to or moderately larger than the incident wavelength. The basic concept was first introduced in 1964 to study the optical properties of molecular aggregates (DeVoe, H., 1964, 1965). Subsequently, prompted by the polarisation of light by interstellar dust and the desire to calculate the approximate extinction, absorption and scattering cross sections for wavelength-sized dielectric grains of arbitrary shape, a comparable scheme was formulated (Purcell, E. M. and Pennymaker, C. R., 1973). Since then, the DDA has undergone major development. A recent erudite review describes this development in detail (Yurkin, M. A. and Hoekstra, A. G., 2007).

There are currently two main publicly available DDA implementations, DDSCAT[3] and ADDA[4], the development of which has provided seminal contributions to the advancement of the DDA. DDSCAT (Draine, B. T. and Flatau, P. J., 2004) is written in FORTRAN 77 and is, at present, the most widely used DDA implementation. It has been developed by Bruce Draine from the Princeton University Observatory and Piotr Flatau from the Scripps Institution of Oceanography, University of California, San Diego. Its development provided several advancements including, among others, an updated expression for the dipole polarisability that incorporated a radiative reaction correction (Draine, B. T., 1988), the application of DFT techniques to the iterative core (Goodman, J. J. et al., 1991), and the introduction of the Lattice Dispersion Relations (LDR) for the prescription of the dipole polarisability (Draine, B. T. and Goodman, J., 1993). A review paper by the authors concisely summarises these contributions (Draine, B. T. and Flatau, P. J., 1994). ADDA (Yurkin, M. A. et al., 2007) is an extremely advanced C code that has been developed (over a period of more than 10 years) by Maxim Yurkin and Alfons Hoekstra at the University of Amsterdam. It was the first parallel implementation of DDA and from its initial incarnation has been able to perform single DDA computations across clusters of autonomous computers (Hoekstra, A. G., 1994; Hoekstra, A. G. and Sloot, P. M. A., 1995; Hoekstra, A. G. et al., 1998). This enables the code to overcome the restrictions on the number of dipoles that can be used due to the memory limitations of a single computer. The code also incorporates several integrated iterative solvers, several prescriptions for prescribing the dipole polarisabilities, and a host of other functionality (Yurkin, M. A. and Hoekstra, A. G., 2008). Other DDA implementations, which are not publicly available, see (Yurkin, M. A. et al., 2007) for further details, include SIRRI (Rahola, J., 1996; Lumme, K. and Rahola, J., 1998); a FORTRAN 90 code that has been developed by Simulintu Ltd in Finland, and ZDD (Zubko, E. et al., 2003); a C++ code that has been developed by Evgenij Zubko and Yuriy Shkuratov from the Institute of Astronomy, Kharkov National University in the Ukraine. For interested parties, a recent paper performs an in-depth comparison of these four codes (Penttila, A. et al., 2007).

Put succinctly, the DDA is an extremely flexible and powerful numerical scheme which replaces a continuum scatterer by an ensemble of $N_d$ dipoles situated on a cubic lattice. Each of the $N_d$ dipoles is driven by the incident electric field and by contributions from the $N_d - 1$ other dipoles. Eq. 1 outlines the kernel of this scheme which is a very large and fully populated set of $3N_d$ complex linear equations.

$$\frac{\mathrm{k}^2}{\epsilon_0} \sum_{j=0}^{N_d-1} \mathbf{A}_{ij}\mathbf{P}_j = \mathbf{E}_{inc,i} \;\; i = 0, \ldots, N_d - 1 \qquad (1)$$
$$\mathbf{A}_{ii} = {\alpha_i}^{-1}$$

[3] http://www.astro.princeton.edu/∼draine
[4] http://www.science.uva.nl/research/scs/Software/adda/index.html

$$\mathbf{A}_{ij} = \frac{e^{\imath \mathrm{k} r_{ij}}}{4\pi {r_{ij}}^3} \left\{ \left[ \frac{\imath \mathrm{k} r_{ij} - 1}{\mathrm{k}^2 {r_{ij}}^2} \right] \left[ 3\mathbf{r}_{ij} \otimes \mathbf{r}_{ij} - {r_{ij}}^2 \mathbf{I} \right] + \left[ \mathbf{r}_{ij} \otimes \mathbf{r}_{ij} - {r_{ij}}^2 \mathbf{I} \right] \right\}$$

Eq. 1 must be solved for the unknown polarisations, $\mathbf{P}_j$, from which the scattering properties can be calculated. Here, k is the wavenumber, $\epsilon_0$ is the permittivity of free space, $\alpha_i$ is the polarisability of dipole $i$, $\mathbf{E}_{inc,i}$ is the incident electric field at dipole $i$, $\mathbf{I}$ is the identity matrix, $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$, $r_{ij} = |\mathbf{r}_{ij}|$ and $\mathbf{r}_{ij} \otimes \mathbf{r}_{ij}$ is the dyadic product, which is a special case of the Kronecker product for vectors of equal order.

$$\mathbf{r}_{ij} \otimes \mathbf{r}_{ij} = \begin{pmatrix} {r_x}^2 & r_x r_y & r_x r_z \\ r_y r_x & {r_y}^2 & r_y r_z \\ r_z r_x & r_z r_y & {r_z}^2 \end{pmatrix} \tag{2}$$

For systems based on three-dimensional lattices, the computational requirements are extremely sensitive to the size of the lattice used. Even a small increase in the dimensions of the lattice significantly increases the computational overheads. As a result, the true potential of the host algorithm is often encumbered by the severe and often prohibitive burdens imposed by the computational scale required to produce meaningful results. In these cases, it is crucial that every possible algorithmic or computational enhancement be exploited in order to augment their applicability. OpenDDA is a novel computational framework for the DDA which has been custom-built from scratch in the C language. The new framework incorporates a variety of algorithmic and computational optimisations, enabling researchers to take full advantage of the different architectures that exist in the current generation of computational resources and beyond. Apart from the core optimisations, the new framework also incorporates an extremely intuitive and user-friendly interface for specifying the simulation parameters, including a sturdy and forgiving parser and a straightforward human-readable input control file.

The paper is organised as follows. Section 2 discusses the application of bit-fielding to streamline the storage of the target occupation and composition data. Section 3 deals with the design and development of a new highly-optimised DFT kernel to calculate the matrix-vector products within the iterative solution schemes. Section 4 discusses the custom-built iterative schemes that have been included to support the new DFT kernel. Section 5 describes the extension of this formulation to both shared-memory and distributed-memory parallel environments. Section 6 provides various various serial and parallel benchmarks to corroborate the assertions and illustrate the implications of the new optimisations. Finally, Section 7 presents concluding remarks.

## 2 Bit-fielding of target occupation and composition data

In many applications, the range provided by the standard integer data types far exceeds the maximum range associated with the task at hand, and as a result, for large datasets, most of the bits are superfluous, constituting a serious waste of resources. Bit-fielding is a way to overcome this issue by storing data within the bits rather than the integers themselves. The entries can be set & tested by taking the appropriately shifted bit-mask and applying the relevant bitwise operator, LOGICAL OR for setting & LOGICAL AND for testing. The simplicity and speed of bitwise operations makes bit-fielding an extremely efficient method of storing non-negative integer based data.

The DDA's theoretical strength, which is also its computational weakness due to imposed computational overheads, is its volumetric nature. Its ability to treat arbitrarily shaped particles with complex internal structures derives from the fact that each constituent site in the computational domain has separate integer flags describing its properties. Rather than storing the spatial extent of the target as the Cartesian coordinates of each lattice site in the computational domain and a

Boolean occupation flag to indicate whether or not the site is occupied, requiring four integers, OpenDDA simply stores a single bit-fielded occupation flag per lattice site. The Cartesian coordinates need not be stored explicitly as, if required, are readily available as the loop control variables or via direct extraction from the array index itself. The optimisation of this integer data storage reduces the memory required for the spatial description of the target by a factor of 128.

The DDA also requires the specification of the material properties of the target on a dipole-by-dipole basis. This is done by associating each complex refractive index value with a distinct integer flag, and for each dipole, storing three integers to define the material properties in the $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ & $\hat{\mathbf{z}}$ directions. To compress this composition data, OpenDDA employs an extension of the bit-fielding theory to non-Boolean data sets. The new framework automatically calculates the requisite bit-mask, i.e., for four disparate materials, $0 \rightarrow 3$, two bits are all that is required, and the bit-mask is $3_{dec} = 11_{bin}$. For this example, the memory required for the integer data associated with the compositional description of the target is reduced by a factor of 16.

To put the optimisations in perspective, for $K = J = P = 500$, $N = K \times J \times P = 125 \times 10^6$, where $K$, $J$ & $P$ are the $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ & $\hat{\mathbf{z}}$ dimensions of the computational domain, respectively, with four materials, the memory requirements for the integer data associated with the spatial and compositional description of the target are reduced from 3.26GB to just over 0.1GB. Furthermore, the computational overhead incurred by the application of this bit-fielding optimisation within OpenDDA is negligible. Firstly, OpenDDA does not require the actual Cartesian coordinates within the iterative core. They are only needed at the initialisation stage for the calculation of the incident electric field at each of the lattice sites. Secondly, due to the nature of the iterative methods and the fact that the iterative solution vectors only include entries corresponding to occupied sites (discussed in §4), the Boolean occupation flags are only tested at the initialisation stage, during the calculation of the dipole polarisabilities, and once at the start and end of the DFT matrix-vector product routine. Although the integer data constitutes only a fraction of the total memory footprint, with most deriving from the requisite complex vectors, these optimisations are not without merit and fundamentally underpin the efficiency of the current distributed-memory implementation of OpenDDA. To facilitate the elimination of zero storage within all iterative solution vectors, each autonomous node requires local knowledge of the entire spatial and compositional specification of the target and thus, unlike the complex vectors within the iterative core, the target specification is not distributed. Bit-fielding allows the full description of the target to be known globally without a prohibitive overhead. With each node having typically a few gigabytes to work with, the original memory requirements associated with the description of the target would have been comparable to the maximum possible allocation per node, negating the possibility of a viable distributed implementation.

# 3 Kernel optimisation

Due to the sheer size of the system produced by Eq. 1, iterative $O(h[g]^2)$ methods, where $g = 3N_d$ and $h$ is the number of iterations required to reach convergence, are employed, which are much faster than direct $O(g^3)$ methods for large $N_d$ and $h \ll N_d$. The cardinal computational hurdle and key to the acceleration of these iterative methods, such as the conjugate gradient method (Barrett, R. et al., 1994; Golub, G. H. and Van Loan, C. F., 1996), are the matrix-vector products required by these schemes. A major breakthrough in this respect was the realisation that by embedding the arbitrarily shaped target in a $N = K \times J \times P$ computational domain and neglecting the diagonal terms, the matrix-vector products can be reformulated as convolutions (Goodman, J. J. et al., 1991). Once the convolutions have been performed, it is then a trivial matter to include the contributions from the previously neglected terms.

The solution method is predicated upon Toeplitz and circulant matrix structure (Davis, P. J., 1994; Golub, G. H. and Van Loan, C. F., 1996). A Toeplitz matrix of order $[n]$, denoted ${}^{[n]}\mathrm{T}$, is

simply a matrix with constant diagonals, whereas a circulant matrix, ${}^{[n]}\mathrm{C}$, is a special case of Toeplitz structure, where the last element of each row is the same as first element of the next row. Any Toeplitz matrix, ${}^{[n]}\mathrm{T}$, can be converted to its circulant counterpart, ${}^{[n']}\mathrm{C}$, where $n' = 2n-1$, by taking the first column, appending the entries of the first row, excluding the initial entry, in reverse order, and completing the diagonals. The key is that any circulant matrix is diagonalised by the corresponding Fourier matrix, ${}^{[n']}\mathrm{F}$ (Van Loan, C. F., 1992), where ${}^{[n']}\mathrm{c}$ is the first column of ${}^{[n']}\mathrm{C}$.

$$
\begin{aligned}
&{}^{[n']}\mathrm{F} = \frac{1}{\sqrt{n'}}\omega^{pq},\ p,q = 0,\ldots,n'-1,\ \omega = e^{-\frac{2\pi\imath}{n'}} \\
&{}^{[n']}\mathrm{F}^{[n']}\mathrm{C}^{[n']}\mathrm{F}^{\mathrm{H}} = diag\left(\sqrt{n'}\left[{}^{[n']}\mathrm{F}\right]{}^{[n']}\mathrm{c}\right) = diag\left(\lambda_C\right)
\end{aligned}
\tag{3}
$$

Using the fact that the Fourier matrix is unitary, ${}^{[n']}\mathrm{F}^{-1} = {}^{[n']}\mathrm{F}^{\mathrm{H}}$, the multiplication of an arbitrary vector with a circulant matrix can be performed as an element-wise multiplication in the Fourier Domain (FD) via the Convolution Theorem, $f \rightleftharpoons \mathrm{F}$, $g \rightleftharpoons \mathrm{G}$, $f * g \rightleftharpoons \mathrm{F.G}$. The symbol '$*$' denotes convolution, '$\rightleftharpoons$' denotes a Fourier transform pair and the symbol '.' implies element-wise multiplication.

$$
{}^{[n']}\mathrm{y} = {}^{[n']}\mathrm{C}^{[n']}\mathrm{x} = {}^{[n']}\mathrm{F}^{-1}diag\left(\lambda_C\right){}^{[n']}\mathrm{F}^{[n']}\mathrm{x} = {}^{[n']}\mathrm{F}^{-1}[\lambda_C].[{}^{[n']}\mathrm{F}^{[n']}\mathrm{x}] \tag{4}
$$

Not only does this reduce the computational complexity of the matrix-vector product from $O(n^2)$ to $O(n' \ln n')$, but since the circulant eigenvalues are a scalar multiple of the DFT of the first column of the circulant matrix, see Eq. 3, the memory is reduced from $O(n^2)$ to $O(n')$. To expedite the calculation of a Toeplitz matrix-vector product, ${}^{[n]}\mathrm{y} = {}^{[n]}\mathrm{T}^{[n]}\mathrm{x}$, the matrix must undergo Toeplitz-to-circulant (T-to-c) conversion, ${}^{[n]}\mathrm{T} \rightarrow {}^{[n']}\mathrm{C}$, and the vector zero-padded to the dimension of the circulant matrix by appending $n-1$ zeroes. Eq. 4 can then be applied. The first $n$ elements of the resultant are the required answer for the original Toeplitz matrix-vector product, with the remaining $n-1$ entries being redundant. Please note that this methodology and the following discussion is only applicable to systems exhibiting Toeplitz structure. Non-Toeplitz structure negates the transformation to circulant structure and hence, the use of DFTs to calculate the matrix-vector products efficiently. In terms of the specific geometry of the target in question, the method is indifferent. For sparse systems, from targets with a low volume filling fraction (VFF), such as dendritic or highly porous structures, the only difference is the amount of memory required by the iterative solution vectors. The solution vectors within the iterative core do not store entries corresponding to vacant lattice sites, a fact that will be discussed in §4. However, the DFT kernel that efficiently calculates the matrix-vector products required by the iterative solvers must account for the entire rectangular computational domain, including the vacant sites, and thus performs the same amount of calculations, irrespective of the VFF.

For the DDA, by embedding the arbitrary target in a rectangular lattice and neglecting the diagonal terms, the requisite matrix-vector products take on the form of a complex-symmetric level-three tensor-Toeplitz matrix of order $[P][J][K]$, denoted ${}^{[P][J][K]}\mathrm{T}$, in which, the constituent tensor elements are complex-symmetric but not necessarily Toeplitz, multiplied by an arbitrary vector of length $N$, denoted ${}^{[N]}\mathrm{x}$, whose constituent elements are Cartesian vectors. Computationally, ${}^{[P][J][K]}\mathrm{T}$ can be considered as an assemblage of nine complex-symmetric level-three single-element Toeplitz matrices, and since the tensor elements are symmetric, only six of the nine tensor components are independent and need to be considered explicitly, $xx, xy, xz, yy, yz, zz \equiv {}^{[P][J][K]}\mathrm{T0}$, ${}^{[P][J][K]}\mathrm{T1}, \ldots, {}^{[P][J][K]}\mathrm{T5}$. Analogously, the arbitrary vector can be split into three separate $x, y, z$ component vectors, ${}^{[N]}\mathrm{x0}$, ${}^{[N]}\mathrm{x1}$, ${}^{[N]}\mathrm{x2}$. Consequently, this means that the components can be treated separately and the full problem is reduced to an ensemble of smaller single-element problems whose contributions will ultimately be manifested through FD products.

At this point, before attempting to decipher the ensuing optimisations of this DFT technique, for readers who are not familiar with the specific algorithmic details associated with the application of DFTs to the solution of such matrix-vector systems, it would be prudent to read Appendix A. This concisely summarises the conventional application of DFTs to these complex-symmetric level-three tensor-Toeplitz systems in the context of the DDA, as prescribed by Goodman, J. J. et al. (1991).

The complete standard algorithm, which encapsulates the conventional implementation, can be summarised as follows:

1. T-to-c conversion and 3D DFT of the six independent tensor components
2. Zero pad and 3D DFT of the three zero-padded vector components
3. FD element-wise multiplication of the tensor and vector components
4. 3D iDFT and normalisation of the components of the resultant
5. Addition of the previously ignored diagonal contributions

Due to the T-to-c conversions that are required for DFT applicability, the six tensor-components arrays contain a high degree of symmetry. Since a $K' \times J' \times P'$ 3D DFT can be carried out as an ensemble of 1D DFTs, $J' \times P'$ in the $\hat{\mathbf{x}}$, $K' \times P'$ in the $\hat{\mathbf{y}}$, and $K' \times J'$ in the $\hat{\mathbf{z}}$ direction, a new algorithm which exploits these symmetries can be constructed. Rather than expanding the six tensor-component arrays from a $K \times J \times P$ Toeplitz structure to a $K' \times J' \times P'$ circulant structure and applying 3D DFTs, as delineated in Fig. 6, the T-to-c conversion and subsequent DFT can be amalgamated into a single computationally efficient operation. Each of the lines in the $\hat{\mathbf{x}}$ direction, ${}^{[K]}\mathrm{t0}_{3\mathrm{D}j,k}, j = 0, \ldots, J-1, k = 0, \ldots, P-1$, in the $\hat{\mathbf{y}}$ direction, ${}^{[J]}\mathrm{t0}_{3\mathrm{D}i,k}, i = 0, \ldots, K-1, k = 0, \ldots, P-1$, and in the $\hat{\mathbf{z}}$ direction, ${}^{[P]}\mathrm{t0}_{3\mathrm{D}i,j}, i = 0, \ldots, K-1, j = 0, \ldots, J-1$, in section '1' in Fig. 6, can be taken in turn and copied into an external 1D array of length $K' = 2K-1$, $J' = 2J-1$, and $P' = 2P-1$, respectively. Each can be treated as the first column of a level-one Toeplitz matrix and the T-to-c conversion performed. The DFTs of these external vectors can be computed, and subsequently, the first $K$, $J$ and $P$ elements of the resultants can be copied back into the 3D tensor-component array, overwriting the original data. This significantly reduces the number of 1D DFTs that have to be performed. Taking into account the zero-padding to ensure that $K'' = K' + kf = 2K - 1 + kf$, $J'' = J' + jf = 2J - 1 + jf$, and $P'' = P' + pf = 2P - 1 + pf$ permits the use of fast DFT algorithms, the number of transforms is reduced by approximately 75% in total:

$$\textbf{x̂ direction, see Fig. 1(a):}\ 6\,[J'' \times P''] \sim \boxed{\mathbf{24}\,[J \times P]} \rightarrow \boxed{\mathbf{6}\,[J \times P]}$$

$$\textbf{ŷ direction, see Fig. 1(b):}\ 6\,[K'' \times P''] \sim \boxed{\mathbf{24}\,[K \times P]} \rightarrow 6\,[(K+kf) \times P] \sim \boxed{\mathbf{6}\,[K \times P]}$$

$$\textbf{ẑ direction, see Fig. 1(c):}\ 6\,[K'' \times J''] \sim \boxed{\mathbf{24}\,[K \times J]} \rightarrow 6\,[(K+kf) \times (J+jf)] \sim \boxed{\mathbf{6}\,[K \times J]}$$

Furthermore, since the full T-to-c conversion, as delineated in Fig. 6, is never explicitly performed, only section '1' in Fig. 6 need be stored, reducing the complex number storage required for the six tensor-component arrays by approximately 88%:

$$6\,[K'' \times J'' \times P''] \sim \boxed{\mathbf{48}\,[K \times J \times P]} \rightarrow 6\,[(K+kf) \times (J+jf) \times (P+pf)] \sim \boxed{\mathbf{6}\,[K \times J \times P]}$$

For the three zero-padded vector-component arrays, the realisation that $\approx \frac{7}{8}$ of each is by very definition, identically zero, can be used to expedite the calculations. This was originally proposed by Hoekstra, A. G. et al. (1998) and has been exploited within ADDA (Yurkin, M. A. and Hoekstra, A. G., 2008). Within OpenDDA, for computational efficiency, the DFTs of the three vector components, the subsequent element-wise FD multiplication of the tensor and vector components, and the iDFTs of the resultant, are amalgamated into a conjoint operation. Unlike the tensor component DFTs,

which are performed in the $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ & $\hat{\mathbf{z}}$ directions, the DFTs of the vector components are performed in the $\hat{\mathbf{y}}$, $\hat{\mathbf{x}}$ & $\hat{\mathbf{z}}$ directions, and the iDFTs are performed in the $\hat{\mathbf{z}}$, $\hat{\mathbf{x}}$ & $\hat{\mathbf{y}}$ directions. The reason for this is twofold. Firstly, it allows the use of a $xz$-scratch plane for the DFTs in the $\hat{\mathbf{x}}$ & $\hat{\mathbf{z}}$ directions, the FD multiplication, and the iDFTs in the $\hat{\mathbf{z}}$ & $\hat{\mathbf{x}}$ directions, which provides significant memory reductions, as only $\approx \frac{1}{4}$ of each of the three full $K'' \times J'' \times P''$ arrays need be stored. Secondly, it supports the distributed FD multiplication in the MPI version of the algorithm (discussed in §5). As a result of this optimisation, the complex number storage for the three vector components, ignoring the storage required for the three $xz$-scratch planes, is reduced by approximately 75%:

$$3\,[K'' \times J'' \times P''] \sim \boxed{\mathbf{24}\,[K \times J \times P]} \rightarrow\ 3\,[K \times J'' \times P] \sim \boxed{\mathbf{6}\,[K \times J \times P]}$$

Each of the three vector components is embedded into a zero-padded $[K \times J'' \times P]$ array, $K$ in the $\hat{\mathbf{x}}$ direction, $J''$ in the $\hat{\mathbf{y}}$ direction, and $P$ in the $\hat{\mathbf{z}}$ direction. Once embedded, the DFT is performed in the $\hat{\mathbf{y}}$ direction. Then, each of the $[K \times P]$ $xz$-planes is, in turn, copied into a $[K'' \times P'']$ external zero-padded scratch $xz$-plane. Subsequently, the DFTs in the $\hat{\mathbf{x}}$ & $\hat{\mathbf{z}}$ directions, the element-wise FD multiplication of the tensor & vector components, and the iDFTs in the $\hat{\mathbf{z}}$ & $\hat{\mathbf{x}}$ directions are performed 'on-the-fly' for this $xz$-plane. The initial $[K \times P]$ section is then copied back into the 3D vector component array, overwriting the original data. Finally, once this has been done for each of the $J''$ $xz$-planes, the iDFT in the $\hat{\mathbf{y}}$ direction is performed.

The reason that this is all possible is due to the fact that only non-zero DFTs, and only iDFTs that will contribute to the final answer, need to be performed. As far as the DFTs are concerned, in the $\hat{\mathbf{y}}$ direction, only the lines ${}^{[J']}\mathrm{x0_{3D}}_{i,k}, i = 0, \ldots, K-1, k = 0, \ldots, P-1$, see Fig. 1 (d), contain any non-zero values and need be considered. Subsequently, in the $\hat{\mathbf{x}}$ direction, only the lines ${}^{[K']}\mathrm{x0_{3D}}_{j,k}, j = 0, \ldots, J'-1, k = 0, \ldots, P-1$, see Fig. 1 (e), need be transformed, and in the $\hat{\mathbf{z}}$ direction, all the lines ${}^{[P']}\mathrm{x0_{3D}}_{i,j}, i = 0, \ldots, K'-1, j = 0, \ldots, J'-1$, see Fig. 1 (f), contain non-zero values and must be included. Taking into account the zero-padding to ensure that $K'' = K' + kf = 2K - 1 + kf$, $J'' = J' + jf = 2J - 1 + jf$, and $P'' = P' + pf = 2P - 1 + pf$ permits the use of fast DFT algorithms, the number of transforms is reduced by approximately 42% in total:

$$\begin{aligned}
&\textbf{$\hat{\mathbf{y}}$ direction, see Fig. 6(d):}\ 3\,[K'' \times P''] \sim \boxed{\mathbf{12}\,[K \times P]} \rightarrow \boxed{3\,[K \times P]} \\
&\textbf{$\hat{\mathbf{x}}$ direction, see Fig. 6(e):}\ 3\,[J'' \times P''] \sim \boxed{\mathbf{12}\,[J \times P]} \rightarrow 3\,[J'' \times P] \sim \boxed{6\,[J \times P]} \\
&\textbf{$\hat{\mathbf{z}}$ direction, see Fig. 6(f):}\ \text{'}\textit{no reduction possible}\text{'}
\end{aligned}$$

The FD element-wise multiplications can be carried out just as detailed in Eqs. 6 → 8 in Appendix A, with one minor caveat. The six independent tensor-component arrays never underwent the full T-to-c conversion, however, since the DFT preserves the mirror symmetries, when accessing Y0 → Y5, the pertinent data can be extracted using the criteria:

$$\begin{aligned}
i &= \begin{cases} i, & \text{if } i < K; \\ K' - i, & \text{if } i \geq K. \end{cases} \quad j = \begin{cases} j, & \text{if } j < J; \\ J' - j, & \text{if } j \geq J. \end{cases} \quad k = \begin{cases} k, & \text{if } k < P; \\ P' - k, & \text{if } k \geq P. \end{cases} \\
&i = 0, \ldots, K'-1 \qquad j = 0, \ldots, J'-1 \qquad k = 0, \ldots, P'-1
\end{aligned}$$

For the iDFT of the resultant of the FD multiplication, i.e., of Q0, Q1 & Q2, see Eqs. 6 → 8 in Appendix A, only transforms that will contribute to the final requisite answer need to be performed. In the $\hat{\mathbf{z}}$ direction, all the lines ${}^{[P']}\mathrm{Q0}_{i,j}, i = 0, \ldots, K'-1, j = 0, \ldots, J'-1$, see Fig. 1 (f), contribute, through the subsequent $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ transforms, to the final answer. Subsequently, in the $\hat{\mathbf{x}}$ direction, only the lines ${}^{[K']}\mathrm{Q0}_{j,k}, j = 0, \ldots, J'-1, k = 0, \ldots, P-1$, see Fig. 1 (e), contribute, through the subsequent $\hat{\mathbf{y}}$ transform, to the final answer, and finally, in the $\hat{\mathbf{y}}$ direction,

only the lines ${}^{[J'']}\mathrm{Q0}_{i,k}, i = 0, \ldots, K-1, k = 0, \ldots, P-1$, see Fig. 1 (d), contribute and need to be considered. Taking into account the zero-padding to ensure that $K'' = K' + kf = 2K - 1 + kf$, $J'' = J' + jf = 2J - 1 + jf$, and $P'' = P' + pf = 2P - 1 + pf$ permits the use of fast iDFT algorithms, the number of transforms is reduced, as with the forward DFTs, by approximately 42% in total:

**$\hat{\mathbf{z}}$ direction, see Fig. 6(f):** '*no reduction possible*'

**$\hat{\mathbf{x}}$ direction, see Fig. 6(e):** $3\,[J'' \times P''] \sim \boxed{12\,[J \times P]} \rightarrow 3\,[J'' \times P] \sim \boxed{6\,[J \times P]}$

**$\hat{\mathbf{y}}$ direction, see Fig. 6(d):** $3\,[K'' \times P''] \sim \boxed{12\,[K \times P]} \rightarrow \boxed{3\,[K \times P]}$

The approximate complexity improvements for both computation and storage are given in Table. 2. Please note that for all three architectural variants, serial, OpenMP (shared-memory), and MPI (distributed-memory), the DFT functionality required by the iterative core is provided by the advanced complex interface of the extremely efficient, highly-portable and well known FFTW package (version 3) (Frigo, M. and Johnson, S. G., 2005).

# 4 Iterative schemes

To avoid unnecessary changes in data format to comply with the requisite structure of some external package, which would only serve to subvert any and all attempts to maximise the performance and streamline the structure of the framework, (like similar DDA implementations) OpenDDA incorporates a relatively large selection of iterative schemes that have been specifically custom-built to integrate seamlessly with the new DFT matrix-vector kernel and support the custom-built domain decomposition and parallel memory allocation schemes in the MPI implementation. Although a variety of different iterative solvers have been employed by Draine, B. T. and Flatau, P. J. (1994); Lumme, K. and Rahola, J. (1994); Flatau, P. J. (1997); Nebeker, B. M. et al. (1998); Rahola J. (1998); Fan, Z. H. et al. (2006); Penttila, A. et al. (2007); Yurkin, M. A. et al. (2007), to date, no definitive "best" algorithm has been identified. To permit further investigation and to maximise the flexibility and usefulness of the end product, OpenDDA includes: Conjugate Gradients (**CG**) (Hestenes, M. R. and Stiefel, E., 1952; Shewchuk, J. R., 1994), Conjugate Gradients Squared (**CGS**) (Sonneveld, P., 1989; Barrett, R. et al., 1994), BiConjugate Gradients (**BiCG**) (Fletcher, R., 1975; Barrett, R. et al., 1994), BiCG for symmetric systems (**BiCG$_{\mathbf{sym}}$**) (Freund, R., 1992), Stabilised version of the BiCG (**BiCGSTAB**) (van der Vorst, H. A., 1992), Restarted, stabilised version of the BiCG (**RBiCGSTAB**) (Sleijpen, G. L. G. and Fokkema, D. R., 1993), Quasi-minimal residual with coupled two-term recurrences (**QMR**) (Freund, R. and Nachtigal, N., 1991, 1992; Bücker, H. M. and Sauren, M., 1996), QMR for symmetric systems (**QMR$_{\mathbf{sym}}$**) (Freund, R., 1992), Transpose-free QMR (**TFQMR**) (Freund, R., 1993) and a variant of BiCGSTAB based on multiple Lanczos starting vectors (**ML(n)BiCGSTAB**) (Yeung, M. and Chan, T. F., 1999).

For computational efficiency, as in ADDA (Yurkin, M. A. et al., 2007; Yurkin, M. A. and Hoekstra, A. G., 2008), the new framework only stores information corresponding to occupied lattice sites, i.e., the elimination of zero storage within all iterative solution vectors. For the majority of cases, this provides a significant reduction in the memory required to store these vectors, as the target being simulated does not fully utilise the computational domain, generally having fewer occupied sites towards the exterior. As a quick example, consider a spherical target based on an $18 \times 18 \times 18$ computational grid. Although the grid contains 5832 lattice sites, for a spherical target, only 3112 are occupied, and dumping the unoccupied data, which is identically zero, equates to a decrease of $\sim 47\%$ in the memory required for the iterative vectors. To facilitate this optimisation, data extension and compression algorithms were created to extend a target vector out to allow the DFT

matrix-vector product, and subsequently, to compress the resultant back into the computationally efficient form with just the entries corresponding to occupied sites.

# 5 Parallel implementations

Over and above the serial improvements, due to the fact that the DFT kernel has been reduced from using 3D to 1D transforms, parallel techniques can be used to further enhance the performance. Full OpenMP-based and MPI-based implementations of the new framework have been created to take proper advantage of both shared-memory and distributed-memory architectures.

For the shared-memory implementation, for thread safety, each thread must have its own DFT scratch arrays. Furthermore, since the aforementioned data-compression algorithm at the end of the DFT kernel operates in-place within the result buffer, to prevent pertinent data loss by multiple processes working within the same array, it has a maximum scalability of three, i.e., one thread per vector component.

For the distributed-memory version, custom 'cyclic-plane' domain decomposition and parallel memory allocation schemes have been developed, which, apart from permitting the efficient location of, and access to, pertinent data, also provide a decrease in memory per node that is directly proportional to the number of available nodes and inherent 'to-the-nearest-plane' load balancing. This is an important point. For the majority of cases, targets do not have a constant number of occupied sites per plane. Since the custom-built iterative solution schemes only store information corresponding to occupied sites, 'cyclic-plane' decomposition provides a much improved data distribution across the available nodes. For example, take the previously discussed spherical target based on an $18\times18\times18$ computational grid, in which only 3112 of the 5832 lattice sites are occupied. If this computational domain is 'block' decomposed across three nodes, then even for this simplistic case, the imbalance of $\sim 47\%$ in the data distribution dwarfs the negligible $\sim 3\%$ imbalance for the 'cyclic-plane' decomposition, see Fig. 2. It is conceded that by 'block' decomposing the constituent planes non-uniformly among the three nodes, so that node '0' is assigned the first seven planes, node '1' is assigned the four middle 4 planes, and node '2' is assigned the remaining seven planes, the load imbalance is reduced to a mere 2%. However, this requires that the domain decomposition and parallel memory allocation algorithms be aware of, and be able to take account of, the user-defined and essentially arbitrary target geometry. In practice, this has not been the case.

To facilitate the distributed DFTs and iDFT required by the new kernel, the new framework incorporates several in-place local and distributed transpose algorithms, which have all been custom-designed to support the 'cyclic-plane' decomposition. They are all capable of treating completely arbitrary systems decomposed across an arbitrary number of nodes. Please note that, unlike the serial and OpenMP versions, which only store section '1' in Fig. 6, for each of the six independent tensor components, the MPI version also stores an extra octant. Rather than storing a $[K \times J \times P]$ array for each of the six components, the MPI version stores a $[K \times J'' \times P]$ array, $K$ in the $\hat{\mathbf{x}}$ direction, $J''$ in the $\hat{\mathbf{y}}$ direction, and $P$ in the $\hat{\mathbf{z}}$ direction. This is to facilitate the distributed element-wise FD multiplication, see Eqs. 6→8 in Appendix A, that is required by the DFT kernel. This extra memory requirement is not an issue as the linear decrease in the memory per node, afforded by the parallel memory allocation, makes the extra storage inconsequential.

# 6 Benchmarking

On top of the architectural specific nature of the implementations, each also exists in float, double and long double precision. For all benchmarks, double precision was used and level three and inter-procedural optimisations were used for the compilations, i.e., '-O3 -ipo' for the Intel compilers and

'-O3 -ipa' for the Pathscale compilers. Furthermore, to standardise the results, all benchmarks were run using the same set of arbitrarily chosen physical parameters, SPHERE, incident wavelength $\lambda = 3.175\,\mu$m, incident polarisation $\hat{\mathrm{e}}_0 = \hat{\mathrm{x}}$, effective radius (radius of a sphere of equal volume) $a = 0.5\,\mu$m, and complex refractive index $m = 1.63631 + 0.372i$[5]. The following list details both the particular hardware and software configurations of each of the benchmarking systems. For each benchmark, the relevant architecture will be referenced simply by the associated tag, FRANKLIN, NEMO or WALTON.

- **FRANKLIN:** Shared-memory, SGI Altix 3700, Red Hat Enterprise Linux AS release 3, IA-64, 24 × Intel Itanium 2, 1.5 GHz, 6 MB cache, 96 GB RAM, Intel Compilers 9.1, FFTW 3.1.2, Intel MKL 8.1.014.
- **NEMO:** Shared-memory, HP xw9300 Workstation, openSUSE 10.2, x86_64, 2 × AMD Dual-Core Opteron 275, 2.2 GHz, 2 MB(2 × 1 MB) cache, 16 GB RAM, Intel Compilers 9.1, FFTW 3.1.2, MPICH2 1.0.5*p*4.
- **WALTON**[6]**:** Distributed-memory, IBM cluster 1350, SUSE LINUX Enterprise Server 9, x86_64, 476 compute nodes, 952 × AMD Opteron Single-Core Opteron 250, 2.4 GHz, 1 MB cache, 412 have 4GB RAM per node, 64 have 8GB RAM per node, QLogic PathScale Compiler Suite, Version 3.0, FFTW 3.1.2, MPICH2 1.0.2.

To affirm the benefits ascribed to the new DFT kernel, the constituent optimisations have been applied sequentially, on FRANKLIN, to show their respective contribution to the cumulative effect. Fig. 3(a) shows the time to compute the matrix-vector product for the original, various hybrid, and new algorithms in a regime that allows comparison with the explicit matrix-vector multiplication. Fig. 3(b) shows long-range timings, where the algorithms have been applied, within their aptitude, to much larger systems. For both the short-range and long-range cases, the timings for the **3D (Orig)** and **1D (Orig)** algorithms are essentially identical. Since the optimisation of the kernel is predicated on the concept of removing superfluous transforms by performing the requisite 3D DFTs as ensembles of 1D transforms, this congruency is crucial as it shows that there is no computational overhead associated with the migration.

For the short-range case, even though the explicit matrix-vector multiplication was performed with the highly optimised Intel MKL (Intel Corporation, 2007), it is obvious that the sheer size of the systems involved negate its use almost immediately. Furthermore, for both plots, the sequential application of the components of the new algorithm is clearly discernible. To put these improvements in perspective, in comparison to the original algorithm, the new algorithm requires approximately 60% less transforms in order to carry out the matrix-vector product, which has a comparably profound influence on the timings.

In the long-range case, the comparison of the original and new algorithms is limited to the nethermost region of the plot. This is due to the prohibitive memory requirements associated with the original algorithm. Note that even for extremely large systems, the new algorithm is notably faster than the original algorithm is for a system that is just a fraction of the size. Both the short-range and long-range cases also include timings for the OpenMP-based kernel across 4, 8 & 16 threads, which unequivocally show the substantial performance gains that can be obtained from shared-memory parallelisation.

An ancillary computational benefit, as shown in Fig. 3(c), is that the new formulation only requires the creation of a few 1D DFT plans, rather than full 3D DFT plans. This is especially

[5]Complex refractive index of ice for $\lambda = 3.175\mu$m from the REFICE routine, originally developed by Warren in 1984 Warren, S. G. (1984), and subsequently improved by Warren, Gao and Wiscombe in 1995 (Unpublished, source code available in the PyARTS package, a python package developed to compliment the Atmospheric Radiative Transfer Simulator - ARTS. Available from the Satellite Atmospheric Science Group, http://www.sat.ltu.se/arts/tools).

[6]ICHEC distributed cluster, Irish Centre for High-End Computing, http://www.ichec.ie

beneficial if using FFTW (Frigo, M. and Johnson, S. G., 1998, 2005) with FFTW_MEASURE to compute the transforms, which instructs FFTW to run and measure the execution time of several DFTs in order to find the best way to compute the transform. As a result, initialisation times become negligible.

In comparison to the original algorithm, the new algorithm requires over 80% less memory to store the tensor and vector components. This is evident from the approximate memory requirement presented in Fig. 3(d), where, for the maximum benchmarked size, corresponding to $K = J = P = 610$, although the original 3D DFT algorithm would require over 257GB, the new algorithm and its OpenMP counterpart only require 52GB to perform the same calculation. This figure also incorporates the memory requirements for the distributed MPI-based implementation, showing the reduction in memory per node that is only limited by the number of nodes available.

To ascertain the performance and scalability of the OpenMP-based shared-memory implementation, a large-scale benchmark was performed across $1, 2, \ldots, 20$ threads on FRANKLIN, see Fig. 4. The system size was chosen to be $K = J = P = 600$, which produced a computational domain with $N = 216 \times 10^6$ sites in total and a representation of the spherical target with $N_d = 113,098,912$ occupied sites. In terms of memory usage, this system required 49GB for 1 thread, with a linear increase of approximately 66MB per thread to approximately 50GB for 20 threads. The extra memory is for the thread private scratch spaces that are required to ensure thread safety. Although this may sound like large amount of memory, in comparison, for the same system, the original algorithm, which is limited to the serial environment, would require approximately 242GB. The time for the matrix-vector product is reduced from 2576 s for 1 thread to just 156 s using 20 threads, a speedup of approximately 16.5. All components show excellent scalability, except of course, the aforementioned data compression algorithm, see §5, which is not an issue as it constitutes only a negligible fraction of the DFT kernel's runtime.

To test the performance and scalability of the MPI-based distributed-memory implementation, the same $K = J = P = 600$ system as for the OpenMP shared-memory benchmark, was run on WALTON, see Fig. 4. Here, due to the fact that each node has only 4GB of physical memory, the benchmark was initialised across 32 nodes and then scaled up to 64 nodes, requiring only 2.2GB per node across 32 nodes, and just 1.2GB per node across 64 nodes. In terms of performance, the timings and scalability are extremely good. The calculation took 108 s across 32 nodes and 55 s across 64 nodes, a speedup of almost 2. To put this in perspective, for the OpenMP-based benchmark, the same system required 2576 s using 1 thread and 156 s using 20 threads. In terms of these results, there are several noteworthy points. Firstly, the total time is quoted as the sum of the computation and communication times. The total communication time includes the time for the forward transpositions of both the six independent tensor components and three zero-padded vector components, and also includes the time for the reverse transposition of the resultant of the FD element-wise multiplication of the components. These transpositions are essentially block events that are controlled by quite complex custom-built transposition algorithms. Even with the prescribed optimisations, due to the enormity of the tensor and vector component arrays, 'in-place' transposition is an absolute necessity. This complication is compounded by the custom-built parallel-memory allocation and 'cyclic-plane' domain decomposition schemes that provide a decrease in memory per node that is directly proportional to the number of active nodes and also provide inherent load balancing. In order to facilitate efficient 'in-place' transposition of the component arrays, the transposition algorithms employ the pre-computation of local 'in-place' transposes in order to simplify and streamline the host algorithm and subsequent data communication. Unfortunately, at present this negates the possibility of overlapping some of the communication with the computation. Secondly, although all communication timings are marginally sub-linear, all computation timings exhibit super-linear speedups. Upon consulting with the ICHEC[7] technical systems team, the peculiarity is believed

[7]Irish Centre for High-End Computing, http://www.ichec.ie

to be an artifact of the initial benchmark. Due to the novel parallel memory allocation algorithms in the distributed-memory implementation of the OpenDDA framework, the memory footprint is decomposed across the active nodes. However, it is believed that the chosen system size and associated memory requirements were fractionally too large, causing slightly degraded performance due to swapping. Since all subsequent speedups are calculated with respect to this initial benchmark, and since the parallel memory allocation and domain decomposition schemes would have solved this allocation/swapping issue as the number of active nodes increased above 32, the subsequent timings would appear super-linear. This is not considered to be a cause for concern as this only introduces a global shift and it is to overall trend that is of importance here.

To test the performance of the full OpenDDA framework, a benchmark was run on NEMO with $K = J = P = 200$, see Table. 3. This produced a computational domain with $N = 8 \times 10^6$ sites in total and a representation of the spherical target with $N_d = 4,188,896$ occupied sites. The convergence tolerance was taken as the default value of $1 \times 10^{-10}$, the maximum number of iterations was set to $1 \times 10^4$, the initial guess was set to zero, and Point-Jacobi preconditioning was used. The performance and scalability of the full OpenDDA framework is extremely good. For this system, it is the simpler schemes that tend to provide the best performance, however, since convergence tends to be slower for larger refractive index values (Draine, B. T., 1988), as the number of iterations increases, the enhanced stability of the more involved iterative methods may prove beneficial, and ultimately, may provide superior convergence. Although it is duly noted that due to the necessity of distributed transposes, the performance of the MPI version will always lag behind that of the OpenMP version, the slight deficit in performance is offset by the decrease in memory that is directly proportional to the number of nodes available.

To put the new custom-built OpenDDA framework in context, several comparative benchmarks were performed between the various architecture specific implementations of OpenDDA, and ADDA; the fastest and most computationally efficient DDA implementation (Penttila, A. et al., 2007). Since the number of iterations to reach convergence is dependent on the user-defined convergence tolerance and on the way in which the stopping criteria are defined within the specific implementation, the comparison focusses, not on the number of iterations, but on the time per iteration. For ADDA, the reason that only the total time for the initialisation of the tensor components is given is that ADDA performs the initialisation of the tensor components as a single conjoint operation, i.e., the computation of the components and the DFTs of these components are perform together. In terms of the quoted memory requirements, for the OpenDDA variants, the approximate memory requirements are obtained from the in-built estimation routines. To obtain an estimate for the approximate memory requirements for ADDA, only 'large' arrays pertaining to the iterative solution have been included. Furthermore, some of the 'large' arrays have been excluded from consideration, as in ADDA, not all relevant 'large' arrays coexist, and thus, some do not contribute to the cumulative memory requirements. For the shared-memory comparison, for $\text{BICG}_{\text{sym}}$, the memory estimations for ADDA include, [*xvec*, *rvec*, *pvec*, *Einc*, *Avecbuffer*, *Dmatrix*, *Xmatrix*, *slices* & *slices_tr*], and for BiCGSTAB & $\text{QMR}_{\text{sym}}$, the estimation also includes [*vec1*, *vec2* & *vec3*]. For the distributed-memory comparison, the estimations also include the contributions from ADDA's MPI communication buffers, [*BT_buffer* & *BT_rbuffer*], and to look at the efficiency of the decomposition of the associated data across the active nodes, apart from the total memory requirements, the results also give the maximum memory requirement per node.

Table 4 shows the results for the shared-memory comparative benchmark run on NEMO using BiCGSTAB, $\text{BICG}_{\text{sym}}$ & $\text{QMR}_{\text{sym}}$, for a spherical target with $K = J = P = 200$, $N = 8 \times 10^6$, and $N_d = 4,188,896$. Please note that although ADDA does employ Intel vectorisation (Bik, A. J. C., 2004) within its basic linear algebra routines, it does incorporate generic OpenMP shared-memory parallelisation. Similarly, Table 6 shows the results for the distributed-memory comparative benchmark run on WALTON using BiCGSTAB, $\text{BICG}_{\text{sym}}$ & $\text{QMR}_{\text{sym}}$, across 4, 8, 16, 32 & 64 nodes, where the relevant system sizes for each of the constituent benchmarks are quoted in Table 5. In

all instances, the new OpenDDA performs extremely favourably. It is worth noting that the custom 'cyclic-plane' decomposition within the OpenDDA framework allows for an extremely even distribution of the data across the active nodes. Even though, for the comparison across 32 nodes in Table 6, OpenDDA requires marginally more memory in total, the maximum memory required by any individual node is actually less.

Finally, as a quick verification of the correct operation of the new OpenDDA framework, a comparison of the extinction, scattering & absorption efficiencies, $Q_{ext}$, $Q_{sca}$ & $Q_{abs}$, for the well proven output of the widely used DDSCAT code and the three architectural variants of OpenDDA was performed on NEMO, see Table 7. All are in good agreement.

# 7 Conclusions

Within the new framework, the bit-fielding of the storage associated with the spatial and compositional specification of the target, not only provides significant reduction in the associated integer storage, but also facilitates the creation of a versatile and efficient distributed-memory implementation.

Although comparable codes, such as ADDA, have pioneered and employ similar optimisations, OpenDDA incorporates a custom-built highly-optimised DFT kernel that performs the requisite 3D DFTs of both the tensor and vector, which are instrumental in the efficient calculation of the matrix-vector products, as ensembles of 1D transforms. By doing so, a significant fraction of the constituent 1D transforms, which are superfluous, can be neglected. In comparison to the "out-of-the-box" algorithm based on 3D DFTs, OpenDDA requires approximately 60% less transforms and over 80% less memory in order to carry out the matrix-vector products. Furthermore, the optimisations permit parallel techniques to be employed to further enhance the performance. Both the resulting OpenMP-based shared-memory and MPI-based distributed-memory implementations possess extremely good scalability and associated performance. On top of these improvements, due to the fact that, computationally, the new framework only requires the creation of a limited number of 1D DFT plans, rather than the full 3D DFT plans required by the original algorithm, the initialisation times, which can be significant, become completely negligible. The new framework also incorporates a relatively large selection of iterative methods that have been custom-built in serial, OpenMP and MPI to integrate seamlessly with the various architecturally specific DFT kernels.

Within the distributed-memory implementation of the new framework, the novel parallel-memory allocation and domain decomposition schemes provide a decrease in memory per node that is directly proportional to the number of nodes and also provide inherent load balancing. Due to the fact that the new framework only stores information corresponding to occupied lattice sites, depending on the target being simulated, there can be quite an imbalance in the number of occupied sites per plane. However, the new framework employs custom-designed 'cyclic-plane' decomposition, and as a result, boasts an extremely well-balanced data distribution across the active nodes.

The framework performs extremely favourably across all architectural variants, with the novel and unique OpenMP-based implementation providing unprecedented parallel performance on shared-memory architectures.

# Acknowledgements

This work was supported by the CosmoGrid Project[8], funded under the Programme for Research in Third Level Institutions (PRTLI) administered by the Irish Higher Education Authority under the National Development Plan and with partial support from the European Regional Development

[8]CosmoGrid Project, http://www.cosmogrid.ie

Fund. In conclusion, the authors would also like to thank the Irish Centre for High-End Computing (ICHEC)[9] and their dedicated Systems Team.

## References

Barrett, R., Berry, M., Chan, T. F., Demmel, J., Donato, J., Dongarra, J., Eijkhout, V., Pozo, R., Romine, C., and van der Vorst, H. (1994). *Templates for the solution of linear systems: Building blocks for iterative methods* (2nd ed.). SIAM, Philadelphia.

Bik, A. J. C. (2004). *The Software Vectorization Handbook: Applying Multimedia Extensions for Maximum Performance*. Intel Press.

Bücker, H. M. and Sauren, M. (1996). A parallel version of the unsymmetric Lanczos algorithm and its application to QMR. Technical Report KFA-ZAM-IB-9606, Central Institute for Applied Mathematics, Research Centre Jülich, Germany.

Davis, P. J. (1994). *Circulant matrices* (2nd ed.). Chelsea Publishing, New York.

DeVoe, H. (1964). Optical properties of molecular aggregates. I. classical model of electronic absorption and refraction. *J. Chem. Phys. 41*, 393–400.

DeVoe, H. (1965). Optical properties of molecular aggregates. II. classical theory of the refraction, absorption and optical activity of solution and crystals. *J. Chem. Phys. 43*, 3199–3208.

Draine, B. T. (1988). The Discrete-Dipole Approximation and its application to interstellar graphite grains. *Ap. J. 333*, 848–872.

Draine, B. T. and Flatau, P. J. (1994). Discrete-Dipole Approximation for scattering calculations. *J. Opt. Soc. Amer. A 11*, 1491–1499.

Draine, B. T. and Flatau, P. J. (2004). User guide for the Discrete Dipole Approximation code DDSCAT 6.1.

Draine, B. T. and Goodman, J. (1993). Beyond Clausius-Mossotti - wave propagation on a polarizable point lattice and the discrete dipole approximation. *Ap. J. 405*, 685–697.

Fan, Z. H., Wang, D. .X, Chen, R. S., and Yung, E. K. N. (2006). The application of iterative solvers in discrete dipole approximation method for computing electromagnetic scattering. *Microwave Opt. Tech. Lett. 48*, 1741–1746.

Flatau, P. J. (1997). Improvements in the discrete-dipole approximation method of computing scattering and absorption. *Opt. Lett. 22*(16), 1205–1207.

Fletcher, R. (1975). Conjugate gradient methods for indefinite systems. In Watson, G. A. (Ed.), *Proceedings of the Dundee Biennial Conference on Numerical Analysis*, pp. 73–89.

Freund, R. (1992). Conjugate gradient-type methods for linear systems with complex symmetric coefficient matrices. *SIAM J. Sci. Comput. 13*(1), 425–448.

Freund, R. (1993). A transpose-free quasi-minimal residual algorithm for non-Hermitian linear systems. *SIAM J. Sci. Comput. 14*(2), 470–482.

[9]Irish Centre for High-End Computing, http://www.ichec.ie

Freund, R. and Nachtigal, N. (1991). A quasi-minimal residual method for non-hermitian linear systems. *Numer. Math. 60*, 315–339.

Freund, R. and Nachtigal, N. (1992). An implementation of the QMR method based on coupled two-term recurrences. *SIAM J. Sci. Comput. 15*(2), 313–337.

Frigo, M. and Johnson, S. G. (1998). FFTW: An adaptive software architecture for the FFT. *Proc. ICASSP 3*, 13811384.

Frigo, M. and Johnson, S. G. (2005). The design and implementation of FFTW3. *Proceedings of the IEEE, (invited paper), Special Issue on Program Generation, Optimization, and Platform Adaptation. 93*(2), 216–231.

Golub, G. H. and Van Loan, C. F. (1996). *Matrix Computations* (3rd ed.). The Johns Hopkins University Press, Baltimore.

Goodman, J. J., Draine, B. T., and Flatau, P. J. (1991). Application of Fast Fourier Transform techniques to the Discrete Dipole Approximation. *Opt. Lett. 16*, 1198–1200.

Hestenes, M. R. and Stiefel, E. (1952). Methods of conjugate gradient for solving linear systems. *J. Res. Nat. Bur. Stand. 49*, 409–436.

Hoekstra, A. G. (1994). *omputer simulations of elastic light scattering*. Ph. D. thesis, University of Amsterdam, Holland.

Hoekstra, A. G., Grimminick, M., and Sloot, P. M. A. (1998). Large scale simulations of elastic light scattering by a fast discrete dipole approximation. *Int. J. Mod. Phys. C 9*, 87–102.

Hoekstra, A. G. and Sloot, P. M. A. (1995). Coupled dipole simulations of elastic light scattering on parallel systems. *Int. J. Mod. Phys. C 6*, 663–679.

Intel Corporation (2007). Intel Math Kernel Library (MKL). Version 9.1.

Lumme, K. and Rahola, J. (1994). Light-scattering by porous dust particles in the discrete dipole approximation. *Astrophys. J. 425*(2), 653–667.

Lumme, K. and Rahola, J. (1998). Comparison of light scattering by stochastically rough spheres, best-fit spheroids and spheres. *J. Quant. Spectrosc. Radiat. 60*(3), 439–450.

Nebeker, B. M., Starr, G. W., and Hirleman, E. (1998). Evaluation of iteration methods used when modelling scattering features on surfaces using the discrete dipole approximation. *J. Quant. Spectrosc. Radiat. Trans. 60*, 493–500.

Penttila, A., Zubko, E., Lumme, K., Muinonen, K., Yurkin, M. A., Draine, B. T., Rahola, J., Hoekstra, H. G., and Shkuratov, Y. (2007). Comparison between discrete dipole implementations and exact techniques. *J. Quant. Spectrosc. Radiat. Trans. 106*(1-3), 417–436.

Purcell, E. M. and Pennymaker, C. R. (1973). Scattering and absorption of light by non-spherical dielectric grains. *Ap. J. 186*, 705–714.

Rahola, J. (1996). *Efficient solution of dense systems of linear equations in electromagnetic scattering calculations*. Ph. D. thesis, Helsinki University of Technology, Finland.

Rahola J. (1998). *Applied Parallel Computing Large Scale Scientific and Industrial Problems: Iterative solution of dense linear systems arising from integral equations*, Volume 1541 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg.

Shewchuk, J. R. (1994). An Introduction to the Conjugate Gradient Method Without the Agonizing Pain.

Sleijpen, G. L. G. and Fokkema, D. R. (1993). BiCGSTAB(L) for linear equations involving unsymmetric matrices with complex spectrum. *Elect. Trans. Numer. Anal. 1*, 11–32.

Sonneveld, P. (1989). CGS, a fast Lanczos-type solver for nonsymmetric linear systems. *SIAM J. Sci. Comput. 10*(1), 36–52.

van der Vorst, H. A. (1992). Bi-CGSTAB : A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. Sci. Comput. 13*(2), 631–644.

Van Loan, C. F. (1992). *Computational frameworks for the Fast Fourier Transform*. Society for Industrial and Applied Mathematics, Philadelphia.

Warren, S. G. (1984). Optical constants of ice from the ultraviolet to the microwave. *Appl. Opt. 23*(8), 218–229.

Yeung, M. and Chan, T. F. (1999). ML(k)BiCGSTAB: A BiCGSTAB variant based on multiple Lanczos starting vectors. *SIAM J. Sci. Comput. 21*(4), 1263–1290.

Yurkin, M. A. and Hoekstra, A. G. (2007). The discrete dipole approximation: An overview and recent developments. *J. Quant. Spectrosc. Radiat. Trans. 106*, 558589.

Yurkin, M. A. and Hoekstra, A. G. (2008). *User Manual for the Discrete Dipole Approximation Code "Amsterdam DDA"*. Version 0.78.2.

Yurkin, M. A., Maltsev, V. P., and Hoekstra, A. G. (2007). The discrete dipole approximation for simulation of light scattering by particles much larger than the wavelength. *J. Quant. Spectrosc. Radiat. Trans. 106*(1-3), 546–557.

Zubko, E., Shkuratov, Y., Hart, M., Eversole, J., and Videen, G. (2003). Backscattering and negative polarization of agglomerate particles. *Opt. Lett. 28*, 1504–1506.

# A Original DFT algorithm

This appendix concisely summarises the conventional application of DFTs to the solution of complex-symmetric level-three tensor-Toeplitz matrix-vector products in the context of the DDA, as prescribed by Goodman, J. J. et al. (1991).

The extension of the level-one formulation described at the beginning of §3 to the treatment of level-three systems, ${}^{[N]}\mathrm{y} = {}^{[P][J][K]}\mathrm{T}\,{}^{[N]}\mathrm{x}$, is obtained by performing the T-to-c conversion on each of the three levels of ${}^{[P][J][K]}\mathrm{T}$ just as it was performed for the level-one system, resulting in a level-three circulant, ${}^{[P'][J'][K']}\mathrm{C}$, where $K' = 2K-1, J' = 2J-1, P' = 2P-1$.

$$
\begin{aligned}
{}^{[P'][J'][K']}\mathrm{C} &= \sum_{k=0}^{P'-1} \left[{}^{[P']}\psi^{k}\right] \otimes \left[{}^{[J'][K']}\mathrm{C}_{k}\right] \\
{}^{[J'][K']}\mathrm{C}_{k} &= \sum_{j=0}^{J'-1} \left[{}^{[J']}\psi^{j}\right] \otimes \left[{}^{[K']}\mathrm{C}_{k,j}\right] \\
{}^{[K']}\mathrm{C}_{k,j} &= \sum_{i=0}^{K'-1} \mathrm{C}_{k,j,i} \left[{}^{[K']}\psi^{i}\right]
\end{aligned}
$$

Here, each of the ${}^{[J'][K']}\mathrm{C}_k$ is a level-two circulant, each of the ${}^{[K']}\mathrm{C}_{k,j}$ is a level-one circulant, the symbol '$\otimes$' denotes the Kronecker tensor product, and ${}^{[n]}\psi$ is the downward-shift permutation matrix of order $[n]$ (Van Loan, C. F., 1992). The arbitrary vector must also be zero-padded to negate the contribution for the circulant extensions on all three levels, ${}^{[N]}\mathrm{x} \rightarrow {}^{[N']}\mathrm{x}$, where $[N'] = [K' \times J' \times P']$. The complimentary Fourier structure which diagonalises this level-three system is acquired from the hierarchical diagonalisation by the Fourier matrices, ${}^{[P']}\mathrm{F}$, ${}^{[J']}\mathrm{F}$ and ${}^{[K']}\mathrm{F}$, i.e., ${}^{[P'][J'][K']}\mathrm{F} = {}^{[P']}\mathrm{F} \otimes {}^{[J']}\mathrm{F} \otimes {}^{[K']}\mathrm{F}$ (Davis, P. J., 1994). As a result, the level-three matrix-vector product, ${}^{[N']}\mathrm{y} = {}^{[P'][J'][K']}\mathrm{C}{}^{[N']}\mathrm{x}$, can be carried out as an element-wise FD multiplication, where ${}^{[N']}\mathrm{c}$ is the first column of ${}^{[P'][J'][K']}\mathrm{C}$.

$$ {}^{[N']}\mathrm{y} = {}^{[P'][J'][K']}\mathrm{F}^{-1} \left[ \sqrt{P'J'K'} \left[ {}^{[P'][J'][K']}\mathrm{F} \right] {}^{[N']}\mathrm{c} \right] . \left[ {}^{[P'][J'][K']}\mathrm{F}{}^{[N']}\mathrm{x} \right] \tag{5} $$

Furthermore, using the realisation that ${}^{[P'][J'][K']}\mathrm{F} \equiv$ 3D DFT allows Eq. 5 to be reformulated in terms of the 3D DFT. To facilitate this, before the aforementioned T-to-c conversions and zero-padding takes place, the first column of each of the six independent tensor component matrices, ${}^{[N]}\mathrm{t0}, \ldots, {}^{[N]}\mathrm{t5}$, and the three vector components, ${}^{[N]}\mathrm{x0}, {}^{[N]}\mathrm{x1}, {}^{[N]}\mathrm{x2}$, must be arranged into separate 3D structures, ${}^{[N]}\mathrm{t0}_{3\mathrm{D}}, \ldots, {}^{[N]}\mathrm{t5}_{3\mathrm{D}}$ and ${}^{[N]}\mathrm{x0}_{3\mathrm{D}}, {}^{[N]}\mathrm{x1}_{3\mathrm{D}}, {}^{[N]}\mathrm{x2}_{3\mathrm{D}}$, respectively. For example, the transformations ${}^{[N]}\mathrm{t0} \rightarrow {}^{[N]}\mathrm{t0}_{3\mathrm{D}}$ and ${}^{[N]}\mathrm{x0} \rightarrow {}^{[N]}\mathrm{x0}_{3\mathrm{D}}$ are governed by ${}^{[N]}\mathrm{t0}_{3\mathrm{D}i,j,k} = {}^{[N]}\mathrm{t0}_{kJK+jK+i}$ and ${}^{[N]}\mathrm{x0}_{3\mathrm{D}i,j,k} = {}^{[N]}\mathrm{x0}_{kJK+jK+i}$, $i = 0, \ldots, K-1, j = 0, \ldots, J-1, k = 0, \ldots, P-1$, respectively.

Once in 3D form, to facilitate the DFT applicability, each of the six tensor-component arrays must undergo T-to-c conversion, e.g., ${}^{[N]}\mathrm{t0}_{3\mathrm{D}} \rightarrow {}^{[N']}\mathrm{c0}_{3\mathrm{D}}$, and the three vector-component arrays must be zero-padded, e.g., ${}^{[N]}\mathrm{x0}_{3\mathrm{D}} \rightarrow {}^{[N']}\mathrm{x0}_{3\mathrm{D}}$. For a level-three system, the T-to-c conversion involves taking each tensor-component array in turn, denoted by section '1' in Fig. 6, treating each line in the $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ & $\hat{\mathbf{z}}$ directions as the first column of a level-one Toeplitz matrix, and mirroring it, excluding the initial entry, across its last entry. The mirrored elements must also by multiplied by the appropriate entry from Table 1. The reason for this is that for a general T-to-c conversion, one would usually append the first row, excluding the initial entry, in reverse order. However, since the nine tensor components are from the dyadic product in Eq. 2, the Toeplitz matrix is either symmetric or non-symmetric by sign only. Hence, it is possible to just mirror the column itself and, if necessary, multiply the mirrored entries by '$-1$'. The $\hat{\mathbf{x}}$ conversion creates section '2', the $\hat{\mathbf{y}}$ conversion creates section '3', and the $\hat{\mathbf{z}}$ conversion creates section '4' in Fig. 6. The zero-padding of the three vector-component arrays is much simpler. Each component is trivially embedded in the $[0 \rightarrow K-1] \times [0 \rightarrow J-1] \times [0 \rightarrow P-1]$ octant of a $K' \times J' \times P'$ array of zeroes, i.e., it would occupy section '1' in Fig. 6, and the rest of the array would be identically zero.

The algorithm requires the DFT of the six tensor-component arrays, ${}^{[N']}\mathrm{c0}_{3\mathrm{D}}, \ldots, {}^{[N']}\mathrm{c5}_{3\mathrm{D}}$, and the three vector-component arrays, ${}^{[N']}\mathrm{x0}_{3\mathrm{D}}, {}^{[N']}\mathrm{x1}_{3\mathrm{D}}, {}^{[N']}\mathrm{x2}_{3\mathrm{D}}$, yielding $\mathrm{Y1}, \mathrm{Y2}, \ldots, \mathrm{Y5}$ and $\mathrm{X0}, \mathrm{X1}, \mathrm{X2}$, respectively. Note that for the practical application of the theory, as with DDSCAT (Draine, B. T. and Flatau, P. J., 2004) and ADDA (Yurkin, M. A. and Hoekstra, A. G., 2008), for maximum performance, if necessary, the transformations described above are all zero-padded to ensure that the final dimensions of the component arrays permit the use of fast DFT algorithms, i.e., for FFTW (Frigo, M. and Johnson, S. G., 1998, 2005), a dimension which satisfies $2^a 3^b 5^c 7^d 11^e 13^f$, where $a, b, c, d$ are arbitrary and $(e+f \leq 1)$. These are illustrated in Fig. 6 as $kf$, $jf$, and $pf$, and increase the dimensions to $K'' = K' + kf = 2K - 1 + kf$, $J'' = J' + jf = 2J - 1 + jf$, and $P'' = P' + pf = 2P - 1 + pf$. Subsequently, the contribution from each of the individual components, to the full circulant matrix-vector product, is obtained via the element-wise tensor-vector product of their transforms.

$$ \mathrm{Q0}_{i,j,k} = \mathrm{Y0}_{i,j,k}\mathrm{X0}_{i,j,k} + \mathrm{Y1}_{i,j,k}\mathrm{X1}_{i,j,k} + \mathrm{Y2}_{i,j,k}\mathrm{X2}_{i,j,k} \tag{6} $$

$$ \mathrm{Q1}_{i,j,k} = \mathrm{Y1}_{i,j,k}\mathrm{X0}_{i,j,k} + \mathrm{Y3}_{i,j,k}\mathrm{X1}_{i,j,k} + \mathrm{Y4}_{i,j,k}\mathrm{X2}_{i,j,k} \tag{7} $$

$$ \mathrm{Q2}_{i,j,k} = \mathrm{Y2}_{i,j,k}\mathrm{X0}_{i,j,k} + \mathrm{Y4}_{i,j,k}\mathrm{X1}_{i,j,k} + \mathrm{Y5}_{i,j,k}\mathrm{X2}_{i,j,k} \tag{8} $$

$$i \quad = \quad 0, \ldots, K' - 1, j = 0, \ldots, J' - 1, k = 0, \ldots, P' - 1$$

The inverse DFT (iDFT) of Q0, Q1 & Q2 provides y0$_{3D}$, y1$_{3D}$ & y2$_{3D}$, which are the resultant of Eq. 5, in 3D form, and the $x, y, z$ components of the original tensor-Toeplitz matrix-vector product, ${}^{[N]}\mathrm{y} = {}^{[P][J][K]}\mathrm{T}^{[N]}\mathrm{x}$, are obtained by extracting the $[0 \rightarrow K-1] \times [0 \rightarrow J-1] \times [0 \rightarrow P-1]$ octant of y0$_{3D}$, y1$_{3D}$ & y2$_{3D}$, respectively, i.e., '1' in Fig. 6. This reduces the complexity from $O([3N_d]^2)$ operations and storage to $O(12N' \ln N')$ and $O(9N')$, respectively.

Table 1: Signs to account for the non-symmetries in sign only for the T-to-c conversions of the six independent tensor components from the dyadic product in Eq. 2.

| | $xx$ | $xy$ | $xz$ | $yy$ | $yz$ | $zz$ |
|---|---|---|---|---|---|---|
| $x$ DFT | +1 | -1 | -1 | +1 | +1 | +1 |
| $y$ DFT | +1 | -1 | +1 | +1 | -1 | +1 |
| $z$ DFT | +1 | +1 | -1 | +1 | -1 | +1 |

Table 2: Complexity estimations for the original and new algorithms.

| **COMPUTATION** | **Original Algorithm** | **New Algorithm** |
|---|---|---|
| Y0→Y5:$\hat{\mathbf{x}}$ | $O(6N'' \ln K'')$ | $O(6JPK'' \ln K'')$ |
| Y0→Y5:$\hat{\mathbf{y}}$ | $O(6N'' \ln J'')$ | $O(6\,[K+kf]\,PJ'' \ln J'')$ |
| Y0→Y5:$\hat{\mathbf{z}}$ | $O(6N'' \ln P'')$ | $O(6\,[K+kf]\,[J+jf]\,P'' \ln P'')$ |
| X0→X2:$\hat{\mathbf{x}}$ | $O(3N'' \ln K'')$ | $O(3J''PK'' \ln K'')$ |
| X0→X2:$\hat{\mathbf{y}}$ | $O(3N'' \ln J'')$ | $O(3KPJ'' \ln J'')$ |
| X0→X2:$\hat{\mathbf{z}}$ | $O(3N'' \ln P'')$ | $O(3N'' \ln P'')$ |
| Q0→Q2:$\hat{\mathbf{x}}$ | $O(3N'' \ln K'')$ | $O(3J''PK'' \ln K'')$ |
| Q0→Q2:$\hat{\mathbf{y}}$ | $O(3N'' \ln J'')$ | $O(3KPJ'' \ln J'')$ |
| Q0→Q2:$\hat{\mathbf{z}}$ | $O(3N'' \ln P'')$ | $O(3N'' \ln P'')$ |

| **STORAGE** | **Original Algorithm** | **New Algorithm** |
|---|---|---|
| Y0→Y5 | $O(6N'')$ | $O(6\,[(K+kf)\times(J+jf)\times(P+pf)])$ |
| X0→X2 | $O(3N'')$ | $O(3\,[KJ''P])$ |

Table 3: Performance of the full serial, OpenMP & MPI implementations of the new framework on NEMO using 4 processors for the parallel versions. **SV:** Number of starting vectors, **NI:** Number of iterations, **CT:** Convergence time, **SU:** Speedup, **Mem:** Memory required in total, **Node:** Memory required per node.

| | | | **SERIAL** | | **OpenMP** | | | **MPI** | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|
| **Scheme** | **SV** | **NI** | **CT** (s) | **Mem** (GB) | **CT** (s) | **SU** | **Mem** (GB) | **CT** (s) | **SU** | **Mem** (GB) | **Node** (GB) |
| bicg | – | 21 | 1439.7 | 3.1 | 404.3 | 3.56 | 3.2 | 604.4 | 2.4 | 3.9 | 1.0 |
| $\text{bicg}_{\text{sym}}$ | – | 21 | 741.5 | 2.6 | 212.7 | 3.5 | 2.6 | 306.2 | 2.4 | 3.3 | 0.8 |
| bicgstab | – | 12 | 846.6 | 3.1 | 234.7 | 3.6 | 3.2 | 353.8 | 2.4 | 3.9 | 1.0 |
| cg | – | 21 | 746.6 | 2.6 | 209.7 | 3.6 | 2.6 | 320.9 | 2.3 | 3.3 | 0.8 |
| cgs | – | 13 | 911.3 | 3.1 | 257.4 | 3.5 | 3.2 | 382.9 | 2.4 | 3.9 | 1.0 |
| mlbicgstab | 1 | 12 | 836.0 | 3.3 | 238.9 | 3.5 | 3.3 | 341.5 | 2.5 | 4.0 | 1.0 |
| mlbicgstab | 2 | 8 | 811.7 | 3.9 | 246.2 | 3.3 | 3.9 | 347.1 | 2.3 | 4.6 | 1.2 |
| mlbicgstab | 3 | 6 | 842.9 | 4.5 | 266.2 | 3.2 | 4.5 | 369.6 | 2.3 | 5.2 | 1.3 |
| qmr | – | 22 | 3704.6 | 3.5 | 1055.1 | 3.5 | 3.5 | 1636.4 | 2.3 | 4.3 | 1.1 |
| $\text{qmr}_{\text{sym}}$ | – | 21 | 759.6 | 3.1 | 216.7 | 3.5 | 3.2 | 320.7 | 2.4 | 3.9 | 1.0 |
| rbicgstab | 1 | 14 | 996.3 | 3.1 | 284.5 | 3.5 | 3.2 | 456.0 | 2.2 | 3.9 | 1.0 |
| rbicgstab | 2 | 7 | 990.5 | 3.5 | 285.2 | 3.5 | 3.5 | 415.5 | 2.4 | 4.3 | 1.1 |
| rbicgstab | 3 | 4 | 851.1 | 3.9 | 242.2 | 3.5 | 3.9 | 357.7 | 2.4 | 4.6 | 1.2 |
| tfqmr | – | 13 | 1850.3 | 3.5 | 524.7 | 3.5 | 3.5 | 795.1 | 2.3 | 4.3 | 1.1 |

| | **cpus** | **Iteration** | **Product** | **Tensor components** | | | **Memory** |
|---|---|---|---|---|---|---|---|
| | | | | **Creation** | **DFT** | **Total** | |
| **BiCGSTAB** | | (s) | (s) | (s) | (s) | (s) | (GB) |
| **ADDA** | **1** | **77.3** | **37.6** | | | **65.5** | **3.7** |
| **openDDA_S** | **1** | **68.3** | **32.7** | **1.8** | **11.6** | **13.4** | **3.1** |
| **openDDA_OMP** | **4** | **19.8** | **9.2** | **0.5** | **3.3** | **3.8** | **3.2** |
| **$\text{BICG}_{\text{sym}}$** | | (s) | (s) | (s) | (s) | (s) | (GB) |
| **ADDA** | **1** | **40.2** | **37.6** | | | **65.3** | **3.1** |
| **openDDA_S** | **1** | **33.9** | **32.4** | **1.8** | **11.6** | **13.3** | **2.6** |
| **openDDA_OMP** | **4** | **10.5** | **9.6** | **0.5** | **3.3** | **3.8** | **2.6** |
| **$\text{QMR}_{\text{sym}}$** | | (s) | (s) | (s) | (s) | (s) | (GB) |
| **ADDA** | **1** | **40.7** | **38.4** | | | **65.5** | **3.7** |
| **openDDA_S** | **1** | **35.1** | **33.1** | **1.8** | **11.6** | **13.4** | **3.1** |
| **openDDA_OMP** | **4** | **10.2** | **9.2** | **0.5** | **3.3** | **3.8** | **3.2** |

Table 4: Comparative benchmark run on NEMO between ADDA 0.76 and the double-precision serial and OpenMP-based shared-memory variants of OpenDDA using BiCGSTAB, $\text{BICG}_{\text{sym}}$ & $\text{QMR}_{\text{sym}}$, for $K = J = P = 200$, which produced a computational domain with $N = 8 \times 10^6$ sites in total and a representation of the spherical target with $N_d = 4,188,896$ occupied sites.

| | SYSTEM SIZES FOR THE COMPARISON | | |
|---|---|---|---|
| **Nodes** | **K=J=P** | **N** | **$N_d$** |
| 4 | 240 | 13,824,000 | 7,238,592 |
| 8 | 310 | 29,791,000 | 15,599,512 |
| 16 | 380 | 54,872,000 | 28,732,192 |
| 32 | 480 | 110,592,000 | 57,904,928 |
| 64 | 600 | 216,000,000 | 113,098,912 |

Table 5: System size, associated size of the computational domain and the number of occupied sites in the computational domain for the MPI-based comparative benchmark with ADDA across on 4, 8, 16, 32 & 64 nodes on WALTON.

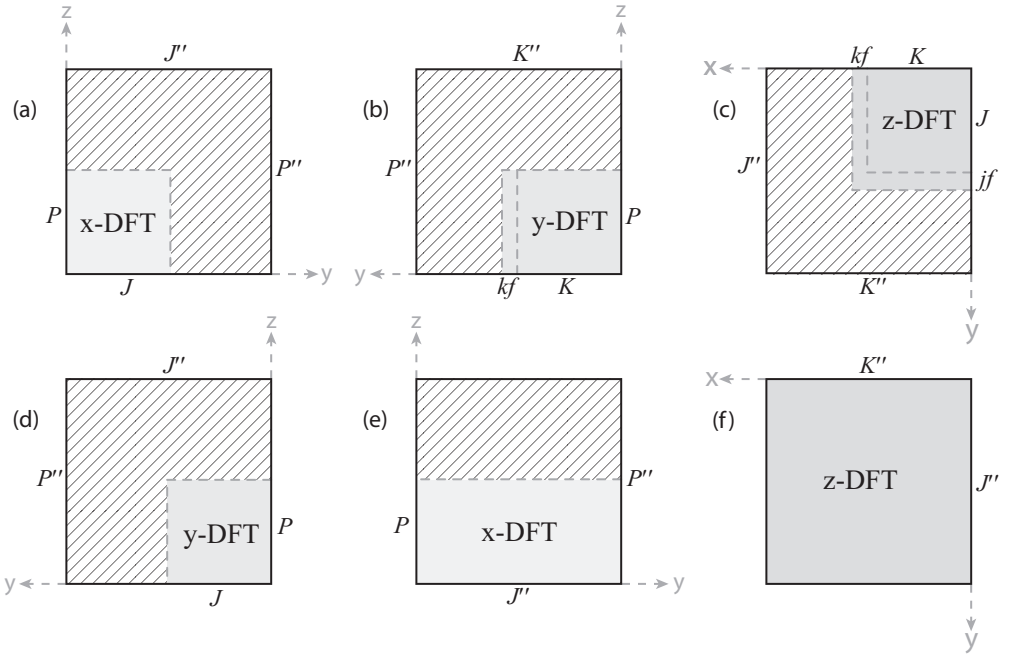


Figure 1: Outlines the optimisation of the requisite 3D DFTs via the implicit T-to-c conversion and simultaneous DFT of the tensor-component arrays and, for the vector components, by only performing transforms that are non-zero or that will contribute to the final answer (a) Tensor components ($\hat{\mathbf{x}}$ direction) (b) Tensor components ($\hat{\mathbf{y}}$ direction) (c) Tensor components ($\hat{\mathbf{z}}$ direction) (d) Vector components ($\hat{\mathbf{y}}$ direction) (e) Vector components ($\hat{\mathbf{x}}$ direction) (f) Vector components ($\hat{\mathbf{z}}$ direction).

| | cpus | Iteration | Product | Tensor components | | | Memory | |
|---|---|---|---|---|---|---|---|---|
| | | | | Creation | DFT | Total | Total | Per node |
| **BiCGSTAB** | | (s) | (s) | (s) | (s) | (s) | (GB) | (GB) |
| ADDA | 4 | 64.8 | 31.2 | | | 41.7 | 7.6 | 2.2 |
| openDDA_MPI | 4 | 54.5 | 25.9 | 17.4 | 0.9 | 18.2 | 6.7 | 1.7 |
| ADDA | 8 | 76.8 | 36.8 | | | 49.9 | 15.4 | 2.3 |
| openDDA_MPI | 8 | 63.9 | 30.6 | 24.4 | 0.9 | 25.3 | 14.7 | 1.8 |
| ADDA | 16 | 79.9 | 36.8 | | | 45.2 | 27.6 | 2.1 |
| openDDA_MPI | 16 | 66.1 | 31.0 | 19.4 | 0.8 | 20.3 | 27.5 | 1.7 |
| ADDA | 32 | 81.5 | 38.2 | | | 44.6 | 54.3 | 2.0 |
| openDDA_MPI | 32 | 71.6 | 34.1 | 23.3 | 0.8 | 24.1 | 54.4 | 1.7 |
| ADDA | 64 | 98.1 | 46.9 | | | 62.9 | 116.7 | 2.2 |
| openDDA_MPI | 64 | 92.2 | 43.7 | 39.4 | 0.8 | 40.2 | 108.9 | 1.7 |
| **$\mathrm{BICG_{sym}}$** | | (s) | (s) | (s) | (s) | (s) | (GB) | (GB) |
| ADDA | 4 | 33.1 | 30.6 | | | 41.2 | 6.7 | 1.8 |
| openDDA_MPI | 4 | 27.3 | 25.9 | 17.3 | 0.9 | 18.2 | 5.7 | 1.4 |
| ADDA | 8 | 37.6 | 35.5 | | | 49.2 | 13.3 | 1.9 |
| openDDA_MPI | 8 | 32.3 | 30.5 | 24.2 | 0.9 | 25.1 | 12.6 | 1.6 |
| ADDA | 16 | 39.9 | 38.2 | | | 45.8 | 23.8 | 1.7 |
| openDDA_MPI | 16 | 30.1 | 30.4 | 19.9 | 0.8 | 20.8 | 23.6 | 1.5 |
| ADDA | 32 | 40.6 | 37.9 | | | 45.7 | 45.6 | 1.7 |
| openDDA_MPI | 32 | 35.8 | 34.1 | 23.0 | 0.8 | 23.8 | 46.6 | 1.5 |
| ADDA | 64 | 50.5 | 47.1 | | | 61.9 | 101.5 | 1.8 |
| openDDA_MPI | 64 | 45.8 | 43.2 | 38.4 | 0.8 | 39.2 | 93.7 | 1.5 |
| **$\mathrm{QMR_{sym}}$** | | (s) | (s) | (s) | (s) | (s) | (GB) | (GB) |
| ADDA | 4 | 33.2 | 32.5 | | | 41.8 | 7.6 | 2.2 |
| openDDA_MPI | 4 | 27.7 | 26.0 | 17.6 | 0.9 | 18.4 | 6.7 | 1.7 |
| ADDA | 8 | 39.2 | 35.9 | | | 48.1 | 15.4 | 2.3 |
| openDDA_MPI | 8 | 32.8 | 30.2 | 24.5 | 0.9 | 25.3 | 14.7 | 1.8 |
| ADDA | 16 | 41.0 | 37.4 | | | 46.3 | 27.6 | 2.1 |
| openDDA_MPI | 16 | 33.3 | 31.2 | 20.1 | 0.8 | 20.9 | 27.5 | 1.7 |
| ADDA | 32 | 40.3 | 38.1 | | | 46.5 | 54.3 | 2.0 |
| openDDA_MPI | 32 | 37.6 | 35.3 | 22.7 | 0.8 | 23.5 | 54.4 | 1.7 |
| ADDA | 64 | 50.8 | 47.5 | | | 61.4 | 116.7 | 2.2 |
| openDDA_MPI | 64 | 46.9 | 43.8 | 38.2 | 0.8 | 39.0 | 108.9 | 1.7 |

Table 6: Comparative benchmark run on WALTON between ADDA 0.76 and the double-precision MPI-based distributed-memory variant of OpenDDA using BiCGSTAB, $\mathrm{BICG_{sym}}$ & $\mathrm{QMR_{sym}}$, across 4, 8, 16, 32 & 64 nodes.

| **DDSCAT** | **OpenDDA (Serial)** | **OpenDDA (OpenMP[4])** | **OpenDDA (MPI[4])** |
|---|---|---|---|
| $Q_{ext}$=1.264 | $Q_{ext}$=1.265 | $Q_{ext}$=1.265 | $Q_{ext}$=1.265 |
| $Q_{abs}$=0.911 | $Q_{abs}$=0.911 | $Q_{abs}$=0.911 | $Q_{abs}$=0.911 |
| $Q_{sca}$=0.353 | $Q_{sca}$=0.354 | $Q_{sca}$=0.354 | $Q_{sca}$=0.354 |

Table 7: Results for the operational verification of the three architectural variants of the new OpenDDA framework. Parameters: CUBE, $K = J = P = 100 \Rightarrow N = N_d = 1 \times 10^6$, normal incidence on one face of the cube, incident wavelength $\lambda = 3.175\,\mu$m, incident polarisation $\hat{e}_0 = \hat{x}$, effective radius $a = 0.5\,\mu$m, and complex refractive index $m = 1.63631 + 0.372i$.

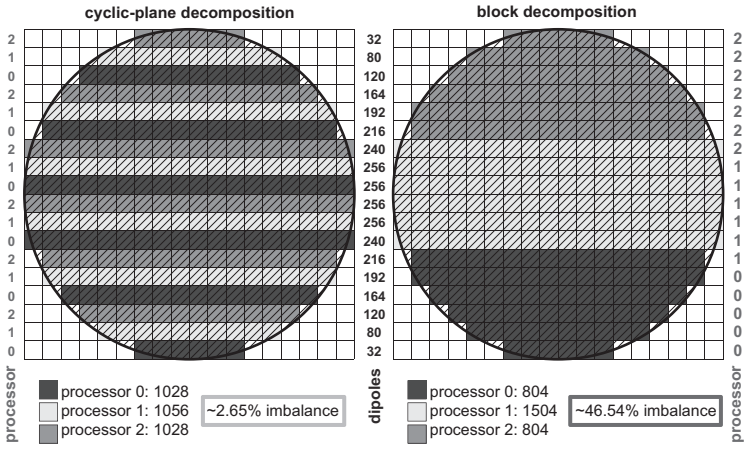


Figure 2: Load balance comparison for block and 'cyclic-plane' decomposition for a simplistic spherical target based on an $18 \times 18 \times 18$ computational grid.

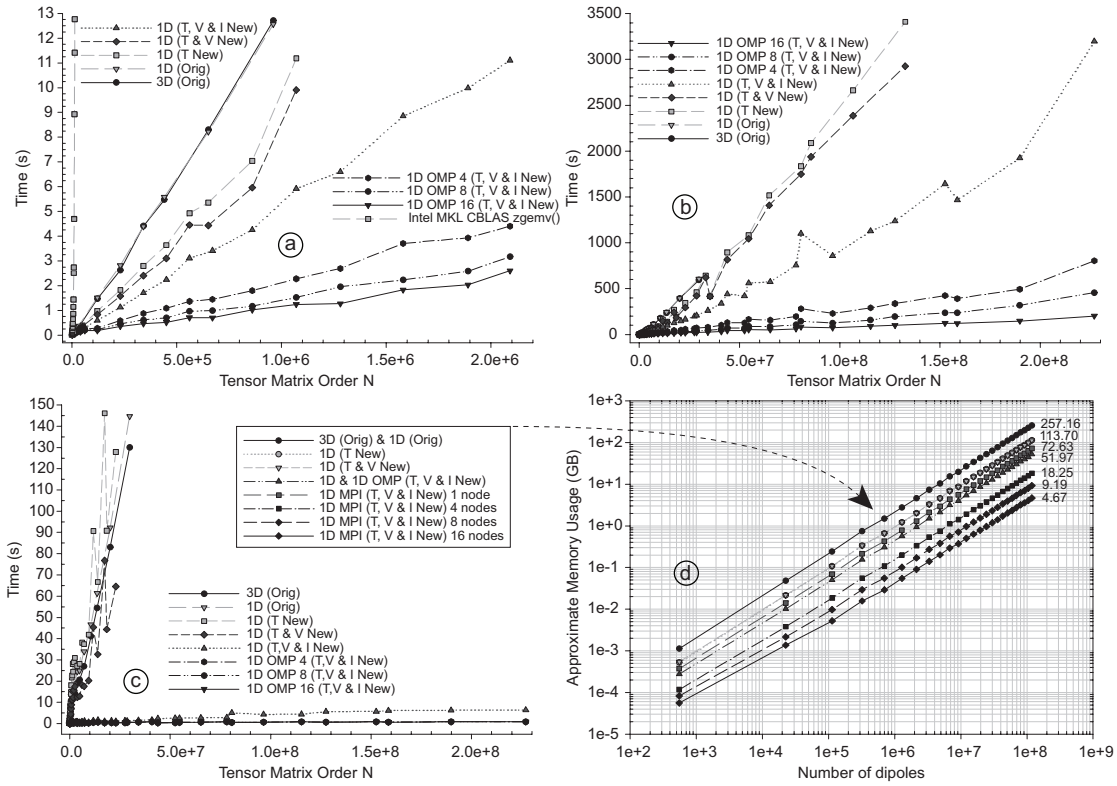


Figure 3: Time and memory required to compute the matrix-vector product on FRANKLIN. *N.B. The label 'Tensor Matrix Order N' refers to a matrix of size N, where the elements are tensors, i.e., the matrix is actually* $3N \times 3N$. (a) Short-range timings; (b) Long-range timings; (c) FFTW DFT plan creation timings; (d) Approximate memory requirements. **3D (Orig):** Original 3D DFT algorithm for all DFTs & iDFTs. **1D (Orig):** Original algorithm using 1D DFTS. **1D (T New):** Hybrid system, new algorithm for the tensor DFTs and the vector DFTs & iDFTs are as in 1D (Orig). **1D (T & V New):** Hybrid system, new algorithm for all tensor & vector DFTS and only the iDFTs are as in 1D (Orig). **1D (T, V & I New):** Complete new algorithm for all DFTs & iDFTs. **1D OMP # (T, V & I New):** New OpenMP algorithm using '#' threads. **1D MPI (T, V & I New) # nodes:** New MPI algorithm using '#' nodes. **Intel MKL CBLAS zgemv:** Explicit calculation using the Intel MKL(Intel Corporation, 2007).

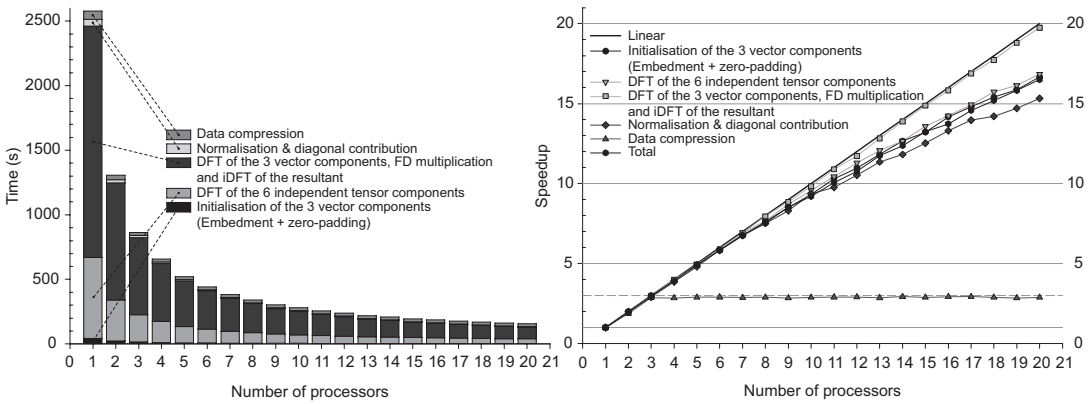


Figure 4: OpenMP scalability and timing run on FRANKLIN for a spherical target with $K = J = P = 600$ and $N_d = 113,089,912$ occupied sites.

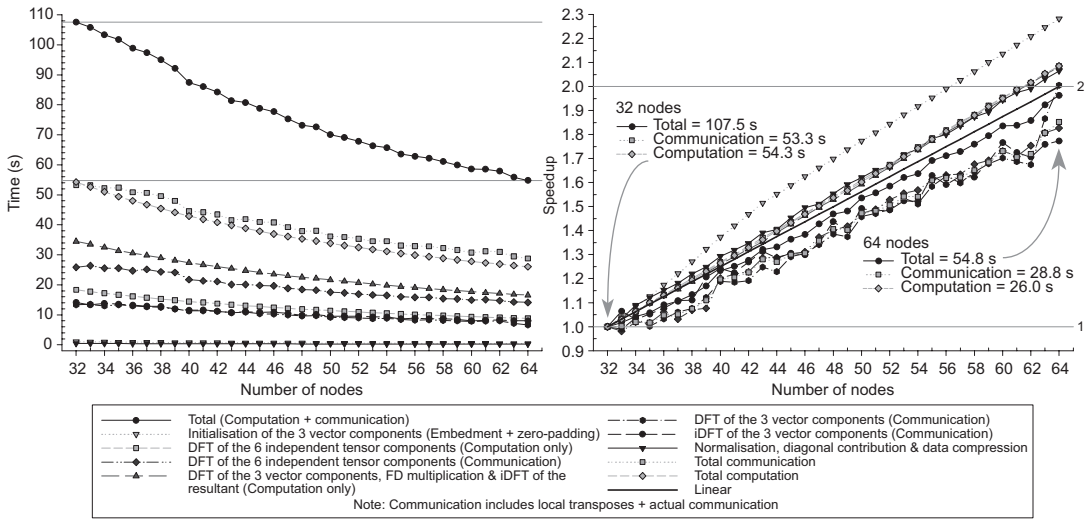


Figure 5: MPI scalability and timing run on WALTON with associated speedups for a spherical target with $K = J = P = 600$ and $N_d = 113,089,912$ occupied sites.

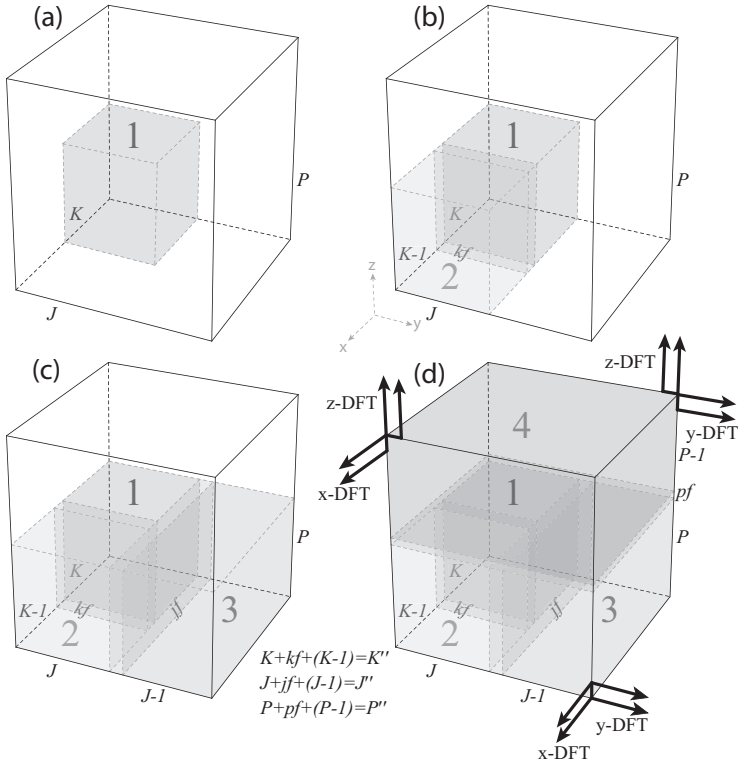


Figure 6: Outlines the T-to-c conversion of the six tensor-component arrays, the zero-padding of the three vector-component arrays and subsequent DFTs & iDFT for the original algorithm.

# Biographies

James Mc Donald received his B.Sc. (1st Class Honours) in Experimental Physics from the National University of Galway (NUI, Galway) in 2002. Subsequently, he worked with the Semiconductor Physics Group in the School of Physics and Astronomy at the University of St Andrews. He returned to Galway to undertake a Ph.D. under the supervision of Dr. Golden and Prof. Jennings at NUI, Galway, which he completed in late 2007. His research interests include high-performance computing and astrophysics (light scattering by non-spherical particles). James Mc Donald gained his Ph.D. in Computational Physics from the National University of Ireland in 2008.

Aaron Golden is College Lecturer in the Department of Information Technology at the National University of Ireland, Galway, and a member of the University's National Centre for Biomedical Engineering Science, and the Centre for Astronomy. He read Natural Sciences at Trinity College Dublin gaining a BA (Mod. Experimental Physics) in 1991, an M.Sc. in Computational Science from the Queen's University of Belfast (1993), and the Ph.D. degree in Astrophysics from the National University of Ireland in 1999. His research interests span topics in astrophysics (magnetospheric emission processes in pulsars and ultracool dwarf stars), computational biology (functional annotation of genomes using artificial intelligence techniques) and high performance computation. This is reflected in the range of journals he has had his work published, from the Astrophysical Journal to Nucleic Acids Research, from Bioinformatics, to Science with 26 peer reviewed published articles and numerous more at conference level. Dr. Golden holds two active Research Frontier Programme Awards from Science Foundation Ireland in Physics and Computer Science, and has supervised 9 Ph.D. students in the 8 years since his appointment at NUI Galway.

S. Gerard Jennings graduated from the National University of Ireland, Galway (NUI, Galway) in 1965 in Experimental Physics with a B.Sc (1st. Class honours), followed by a M.Sc. in 1967. He obtained a Ph.D. in Atmospheric Physics from the University of Manchester in 1971. In 1980, he was appointed Statutory Lecturer in the Department of Experimental Physics, NUI, Galway. He was awarded a D.Sc. from the University of Manchester, 1987. Professor of Experimental Physics at NUI, Galway (1990-present). He was awarded Doctor Honoris Causa by the University of Göteborg in 2000. He was Head of Department of Experimental Physics (1992-1995; 2002-2005) at NUI, Galway, and is currently the director of the Environmental Change Institute (ECI) at NUI, Galway. He is the director of the Atmospheric Research Group, leader of the Climate Change Cluster within the ECI, and management (Secretary) of the Mace Head Atmospheric Research Station, near Carna, Co. Galway. He is Chair of the UK Natural Environmental Research Council (NERC) Polluted Troposphere Scientific Steering Committee (2002-present), a member of the Scientific Advisory Group of the World Meteorological Organisation, Global Atmosphere Watch, (GAW) Aerosol (1997-present), a member of IGACs Coordinating Committee for Aerosol Characterisation and Process Studies Activities, Associate Editor of Journal of Geophysical Research (2002-2005), and a member of the Scientific Steering Committee of EUROTRAC (1992-1999). His research interests include atmospheric physics including aerosol physics, climate change, effects of aerosols on climate, atmospheric field studies at the Mace Head Atmospheric Research Station, environmental pollution, laser propagation in aerosol and cloud media and satellite remote sensing. He has 91 peer refereed journal publications, 1 book, 6 book chapters and 100 conference papers. He has supervised 10 Ph.D. and 14 research M.Sc. students.